

Cooling System Overview: Summit Supercomputer

David Grant, PE, CEM, DCEP
HPC Mechanical Engineer
Oak Ridge National Laboratory
Corresponding Member of ASHRAE TC 9.9
Infrastructure Co-Lead - EEHPCWG

ORNL is managed by UT-Battelle, LLC for the US Department of Energy



U.S. DEPARTMENT OF
ENERGY

Today's Presentation

- System Description
- Cooling System Components
- Cooling System Performance

NOVEMBER 2018 #1


	Rmax	Rpeak	Power
Cores	(TFlop/s)	(TFlop/s)	(kW)
2,397,824	143,500.0	200,794.9	9,783

JUNE 2018 #1

	Rmax	Rpeak	Power
Cores	(TFlop/s)	(TFlop/s)	(kW)
2,282,544	122,300.0	187,659.3	8,806




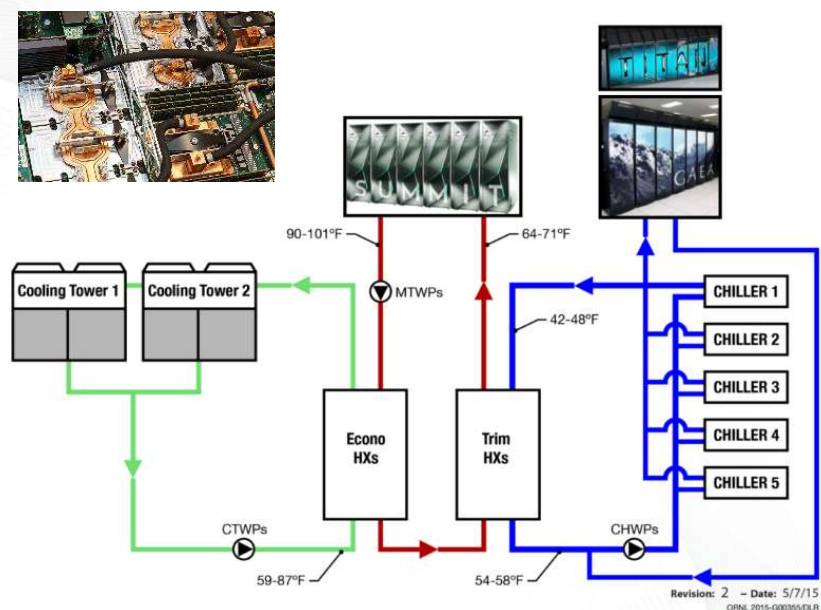
Summit Node Overview



Feature	Titan	Summit
Application Performance	Baseline	5-10x Titan
Number of Nodes	18,688	4,608
Node performance	1.4 TF	42 TF
Memory per Node	32 GB DDR3 + 6 GB GDDR5	512 GB DDR4 + 96 GB HBM2
NV memory per Node	0	1600 GB
Total System Memory	710 TB	>10 PB DDR4 + HBM2 + Non-volatile
System Interconnect	Gemini (6.4 GB/s)	Dual Rail EDR-IB (25 GB/s)
Interconnect Topology	3D Torus	Non-blocking Fat Tree
Bi-Section Bandwidth	15.6 TB/s	115.2 TB/s
Processors	1 AMD Opteron™ 1 NVIDIA Kepler™	2 IBM POWER9™ 6 NVIDIA Volta™
File System	32 PB, 1 TB/s, Lustre®	250 PB, 2.5 TB/s, GPFS™
Peak Power Consumption	9 MW	13 MW

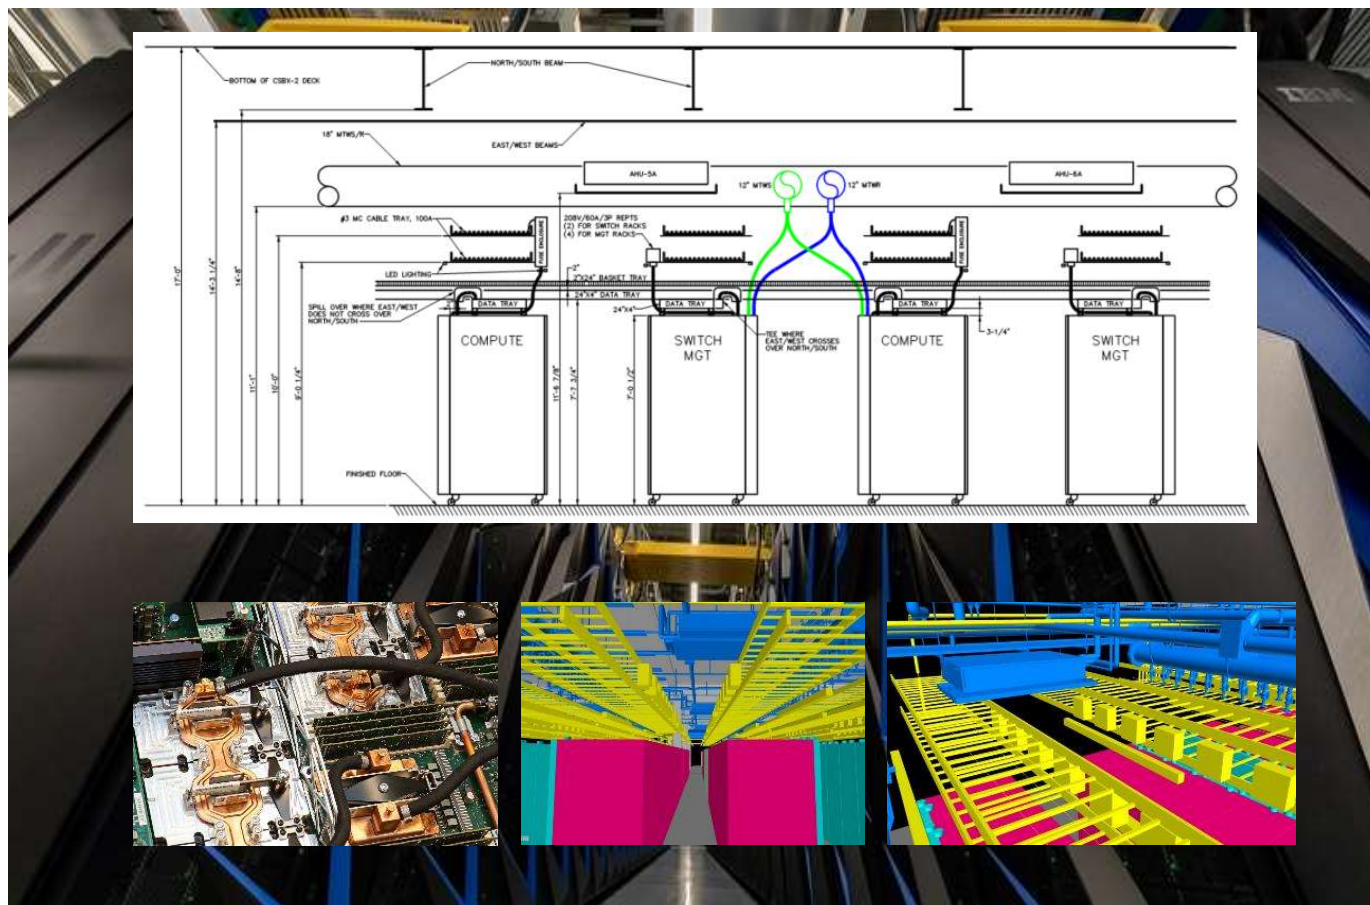
System Description – How do we cool it?

- >100,000 liquid connections




System Description – What's in the data center?

- **Passive RDHXs**- 215,150ft² (19,988m²) of total heat exchange surface (>20X the area of the data center)
 - With a 70°F (21.1 °C) entering water temperature, the room averages ~73°F (22.8°C) with ~3.5MW load and ~75.5°F (23.9°C) with ~10MW load. Note that only ~25% of compute rack load is on the RDHXs.
 - The racks turn over the data center's air volume 2-3 times each minute when under load.
- **CPU cold plates** – 4,105ft² (381m²) of total heat exchange surface
- **GPU cold plates** – 4,448ft² (413m²) of total heat exchange surface
- **Other** - electrical transmission losses, lights, return water piping losses, building envelope, rack radiant, back of rack air exfiltration, VRF AHU fans



System Description – What does the cooling system see in the data center?




System Components – Facility Improvements




PUE – ORNL's Measurement Overview


$$\text{PUE} = \frac{\text{MSBs}(x5) + \text{ATS} + \text{UPSs}(x2) + \text{MSBs}(x2) + \text{Chilled Water Cooling Load} \times \text{Chiller Plant Efficiency}}{\text{MSBs}(x5) + \text{ATS} + \text{UPSs}(x2)}$$




Cooling System Performance - PUE



Cooling System Performance – HPL Runs




Cooling System Performance – Cooling Source






Cooling System Performance – Chilled Water Use Efficiencies

- Chilled Water –
~0.8kW/ton
- Cooling Tower –
~0.2kW/ton



kWh of cooling from
March to October:
22% Chilled Water
78% Cooling Towers





Cooling System Performance – Within the Data Center



Cooling System Performance – Within the Data Center



Thank You



<https://www.olcf.ornl.gov/summit/>

